NIST National Institute of Standards and Technology

INFORMATION TECHNOLOGY LABORATORY

Trustworthy Networking Program

Evaluation of BGP Anomaly Detection and Robustness Algorithms

Kotikapaludi Sriram, Doug Montgomery, Oliver Borchert, Okhee Kim, and Patrick Gleichmann

National Institute of Standards and Technology

(Contact: ksriram@nist.gov; dougm@nist.gov)

NANOG-43, June 2008

This research was supported by the Department of Homeland Security under the Secure Protocols for the Routing Infrastructure (SPRI) program and the NIST Information Technology Laboratory Trustworthy Networking Program.

Outline of the Talk

- **Known / New BGP robustness schemes**
- **Evaluation of BGP robustness algorithms**
  - **Quantitative / comparative analysis of utility**
  - **Preliminary quantitative results**
- **Conclusions / Future Work**

"Blueprint" / Nemecis: Registry Based Algorithm

- For (p, Origin AS) pair from an update:
  - Check for existence of prefix, autnum, and route objects in RIR/IRR
  - Check for consistency between these declared objects by matching OrgID, maintainer, email handle, DNS server, etc.
  - Generate alerts if these checks fail -- full / partial consistency checks

G. Siganos and M. Faloutsos, "A Blueprint for Improving the Robustness of Internet Routing," 2005. http://www.cs.ucr.edu/%7Esiganos/papers/security06.pdf

G. Siganos and M. Faloutsos, "Analyzing BGP policies: methodology and tool," IEEE Infocom, 2004.

PHAS: Prefix Hijack Alert System

- Provide alert messages if:
  - Origin AS set changes
  - New subprefix is added to observed set of subprefixes
  - Last-hop AS set changes

Mohit Lad, Dan Massey, Yiguo Wu, Beichuan Zhang and Lixia Zhang, *PHAS: A prefix hijack alert system*, North American Network Operators Group Meeting (NANOG-38), October, 2006. http://www.nanog.org/mtg-0610/presenter-pdfs/massey.pdf

Mohit Lad, Dan Massey, Dan Pei, Yiguo Wu, Beichuan Zhang and Lixia Zhang, *PHAS: A prefix hijack alert system*, in Proceedings of 15th USENIX Security Symposium (USENIX Security 2006). http://www.cs.ucla.edu/~mohit/cameraReady/ladSecurity06.pdf

PGBGP: Pretty Good BGP

Old Version of the Algorithm

- Observed and "unsuspicious" (prefix, Origin AS) pairs based on update history and RIB entries over the last *h* days (*h* = 10 days) are recorded
- The anomaly detector also eliminates old routes (older than 10 days) if they are no longer active
- A new update is considered suspicious if the origin AS is not in the history record; the update is propagated with lower local pref
- A subprefix is always considered suspicious and quarantined
- The quarantine lasts for suspicious period of *s* hours (*s* = 24 hours); if subprefix is not withdrawn during that time, then the update is propagated

One Weakness of Old PGBGP

From NANOG discussions back in 2006

Q: Panix's first, obvious countermeasure aimed at restoring their connectivity -- announcing their own address space split in half -- would *also* have been considered suspicious, since it gave two "sub-prefixes" of what ConEd was hijacking?

A: [Here] things get a little more subtle. We have considered allowing the trusted originator of a prefix to split the space among itself and those downstream of it without considering that suspicious behavior.

Note: This was part of the Q&A after the paper on PGBGP was presented by J. Karlin at NANOG-37. http://www.nanog.org/mtg-0606/pdf/josh-karlin.pdf

New Version of PGBGP

- From an updated new version of PGBGP paper:
  - "PGBGP would not interfere if an AS announces sub-prefixes of its own prefixes in order to gain traffic back during a prefix hijack."

Josh Karlin, Stephanie Forrest, and Jennifer Rexford, "Pretty Good BGP: Improving BGP by Cautiously Adopting Routes," The 14th IEEE International Conference on Network Protocols, November 2006. http://www.cs.unm.edu/~treport/tr/06-06/pgbgp3.pdf

Potential Weaknesses of (New) PGBGP

- The short-span historical view (last ten days) has the following negative implications:
  - PGBGP will typically unnecessarily lower local-pref on path announcements due to multi-homing related AS origin change.
  - If a malicious user observes a prefix withdrawal by genuine origin AS and announces the prefix at that time, the malicious path propagates with a lower local-pref value and will be used *(*Effectively - *False Negative).*
  - If the prefix owner sometimes announces sub-prefixes in conjunction with multi-homing related AS origin change, PGBGP will quarantine the announcements.

New Integrated Approach

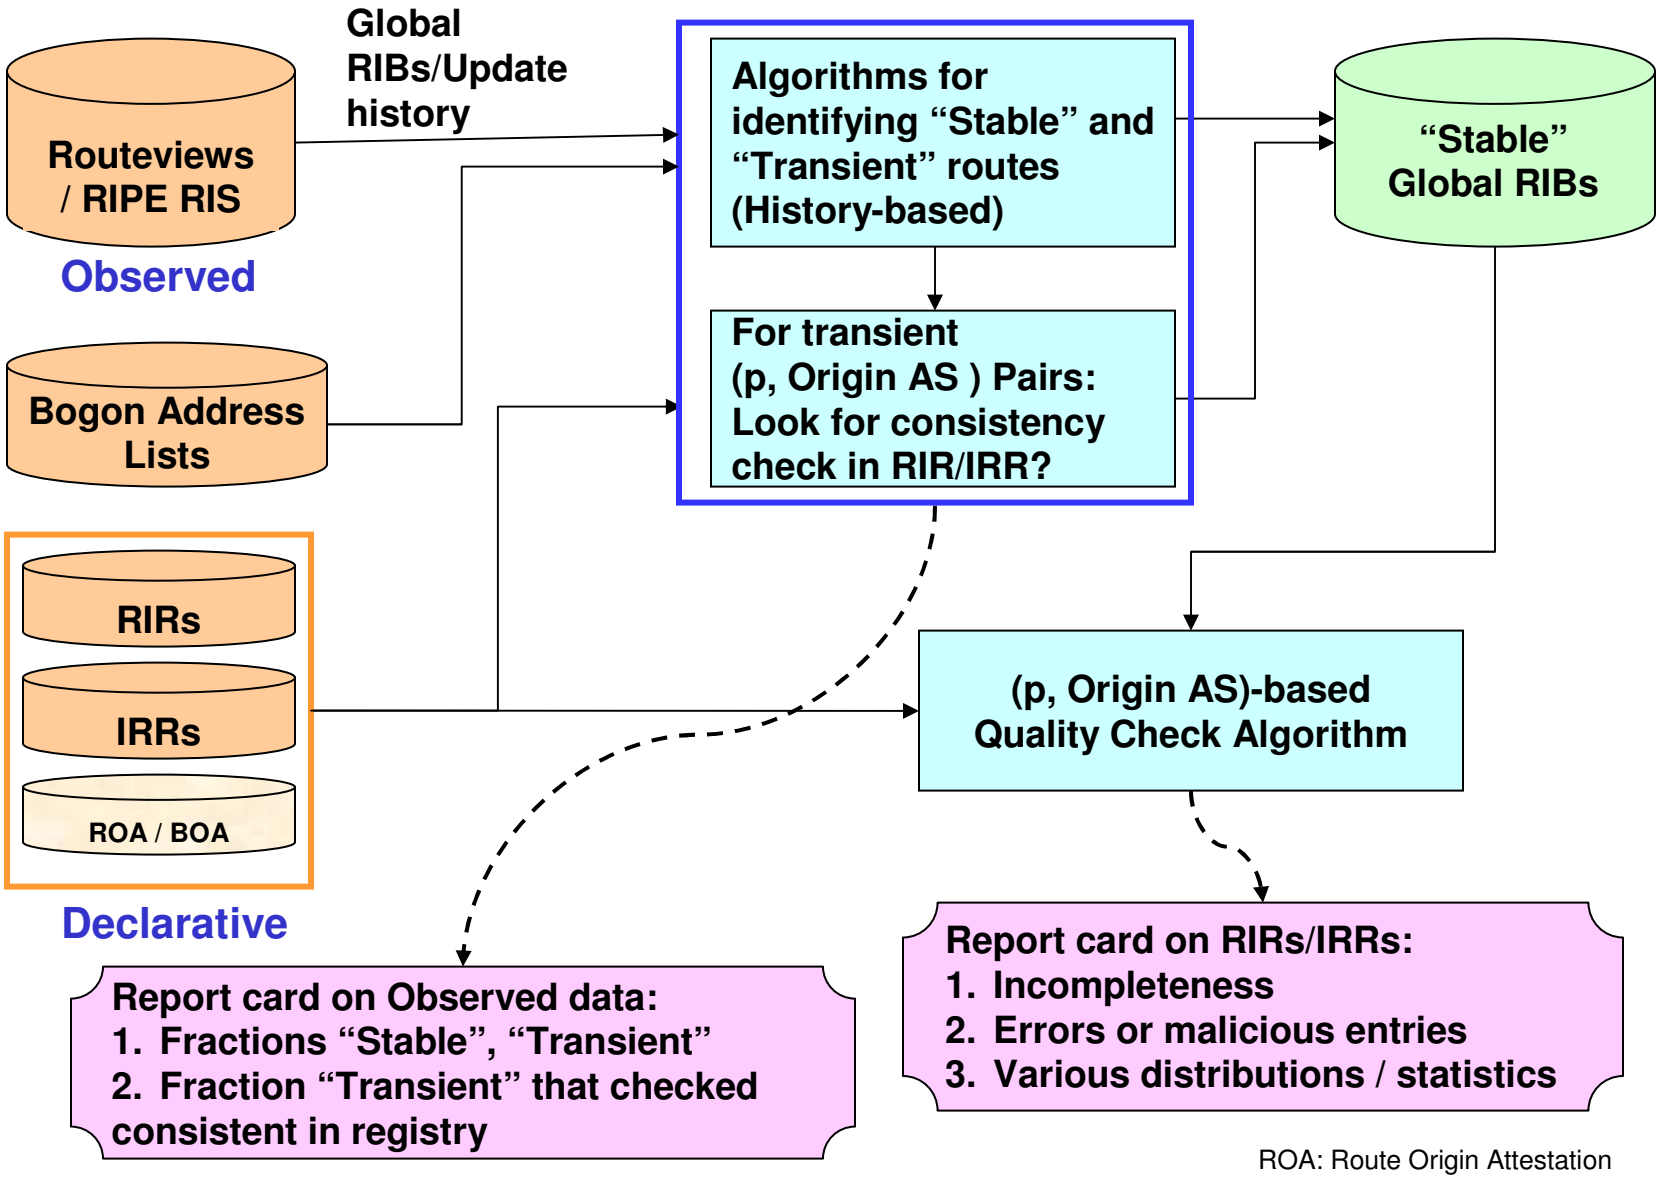

ROA: Route Origin Attestation

BOA: Bogon Origin Attestation

Registry Based Algorithm

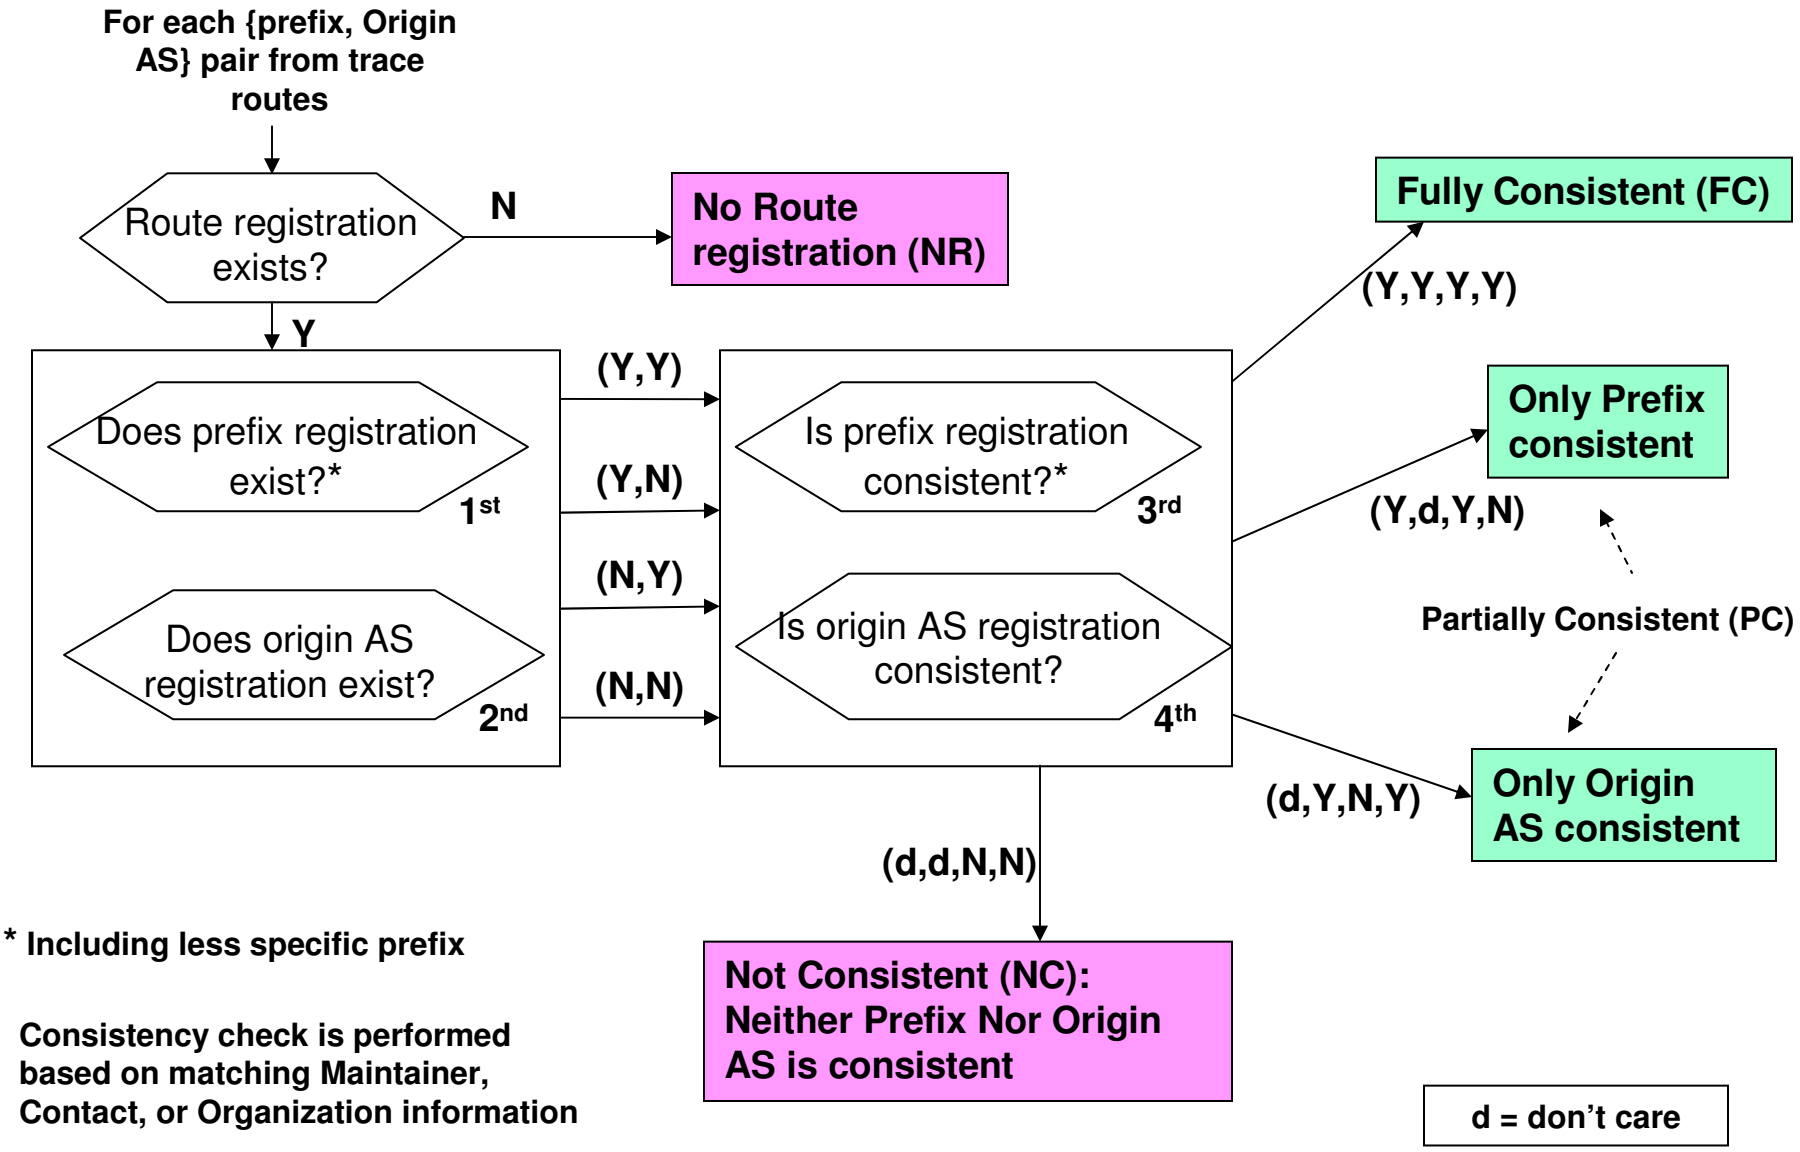

Algorithm for Determining Stability of (p, Origin AS) in the Trace Data

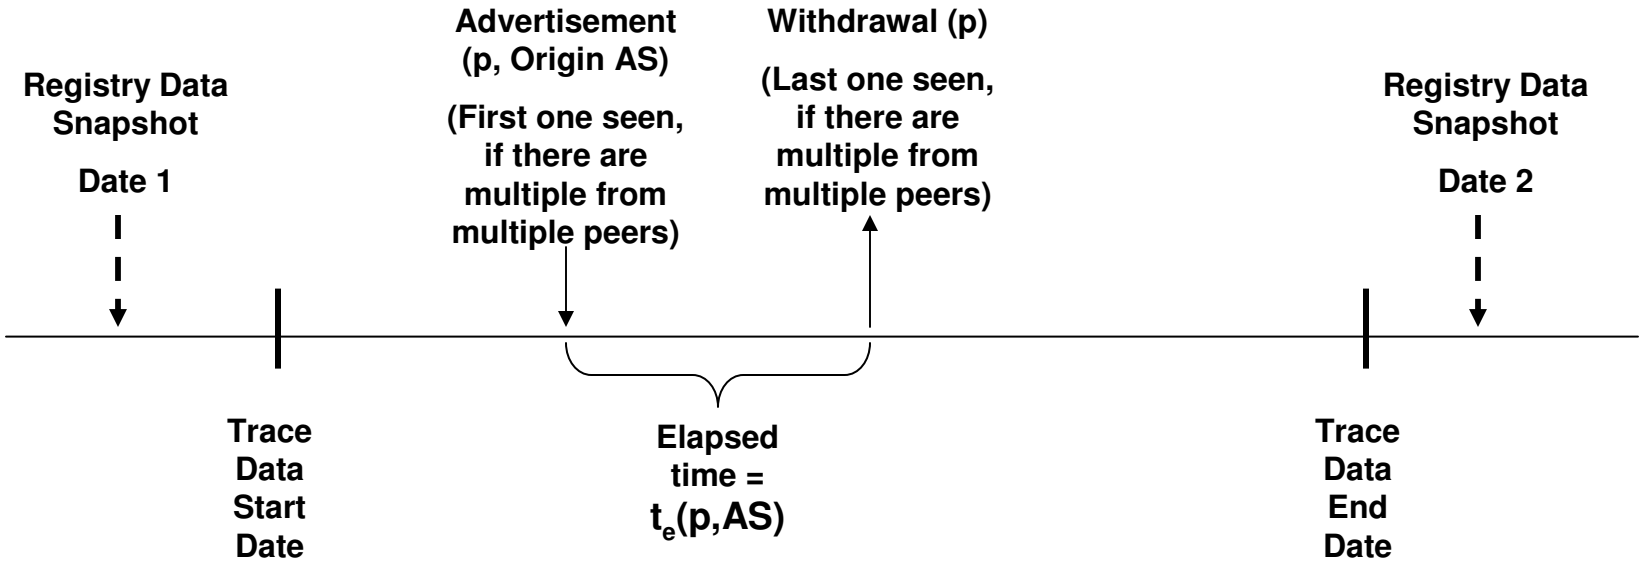

- If (p, AS) had no withdrawal after the advertisement, set t_e(p, AS) = 10^6 hours
- If t_e(p, AS) $\geq$ 48 hours, then (p, AS) is a stable (prefix, Origin AS) pair
- If t_e(p, AS) < 48 hours, then (p, AS) is an unstable (prefix, Origin AS) pair
- Update data is initialized with stable (i.e., at least 48 hours) RIB entries
- Compare each of the two snap shots of the registry data with the stable/unstable sets of historical (prefix, Origin AS) to corroborate

Outline of the Talk

- Known / New BGP robustness schemes
- Evaluation of BGP robustness algorithms
  - Quantitative / comparative analysis of utility
  - Preliminary quantitative results
- Conclusions / Future Work

Origin AS Approval Check List: Comparison

		Which checks are included in each approach?			
	Checks/Questions	Registry-based approach	Trace-data based approach (PGBGP)	Simple Hybrid	Enhanced Hybrid
Q1.	Is prefix registered (same or less specific)?	√		√	√
Q2.	Is there a route registered (with same or less specific prefix and origin AS)?	√		√	√
Q3.	Is announced (p, origin AS) fully consistent with corresponding registry objects in RIR/IRR?	√		√	√
Q4.	Is announced (p, origin AS) partially consistent with corresponding registry objects RIR/IRR?	√		√	√
Q5.	Was (p, origin AS) seen in RIB in the last *h* (= 10) days? (Also, if it was suspicious, did it remain in RIB beyond the suspicious period of *s* (= 24) hours?)		√	√	
Q6.	Would a less specific prefix with the same origin AS pass the test in Q5?		√	√	
Q7.	Was prefix previously announced by the same origin AS and remained stably (48 hrs or more) in the RIB over the observation period (*d* months)?				√
Q8.	Would a less specific prefix with the same origin AS pass the test in Q7?				√
Q9.	Is the peering rank of the origin AS high or medium?				√

Algorithm Robustness Checklist

	Algorithmic Features	Registry-based approach	Trace-data based approach (PGBGP)	Simple Hybrid	Enhanced Hybrid
Data Sets	Utilization of self-consistent registry objects	Yes	No	Yes	Yes
	Utilization of update history	No	Yes	Yes	Yes
	Utilization of historical RIB entries	No	Yes	Yes	Yes
Situations Handled	Pass a subprefix announcement if a less specific prefix with same origin AS could be passed	Yes	Yes	Yes	Yes
	False Positives: Alert raised when genuine prefix owner announces multi-homing related AS origin change	Moderate probability	High probability	High probability	Low probability
	Alert raised when attacker announces a prefix after sensing it has just been withdrawn	Yes	No	Path propagates (lower pref)	Yes
	Pass a subprefix announcement in conjunction with multi-homing related AS origin change	Moderate probability	Low probability	Low probability	High probability

*** This is a ballpark qualitative assessment; subject to corroboration using extensive quantitative studies.**

Comparative Analysis of Existing and Enhanced Algorithms

- We have encoded Registry-based, Trace-data-based and Enhanced Hybrid algorithms for evaluation
- Trace-data based algorithm is a variant from PGBGP (see slides 11, 30, 31)
- Algorithms are run on top of the NIST TERRAIN framework
  - Unified database of Registry / Trace data (RIRs, IRRs, RIPE-RIS, Routeviews)
- Tested and compared the algorithms

Comparative Analysis of Existing and Enhanced Algorithms (Contd.)

For the purpose of this presentation:

- Results focus on Origin AS validation
- Focus on RIPE RIR/IRR and RIPE RIS data
  - (Prefix, Origin AS) pairs are filtered based on RIPE NCC addresses
- Six month trace-data (January through June 2007); initialized with stable (i.e., at least 48 hours) RIB entries
- Registry data – just before and after the six month window
- Preliminary comparison results follow

Some Caveats Apply

- This presentation is mainly to demonstrate the capability and to solicit feedback on approach
- Quantitative results are subject to change when the following enhancements to the study are made (ongoing work)
  - Consideration of registry data from all regions
  - Reconciling related registry objects in different regional registries
  - Consideration of multiple trace-data collectors (here we considered trace-data from RRC02 only)

Classification of Observed (p, OAS) Pairs According to Stability / Consistency Scores

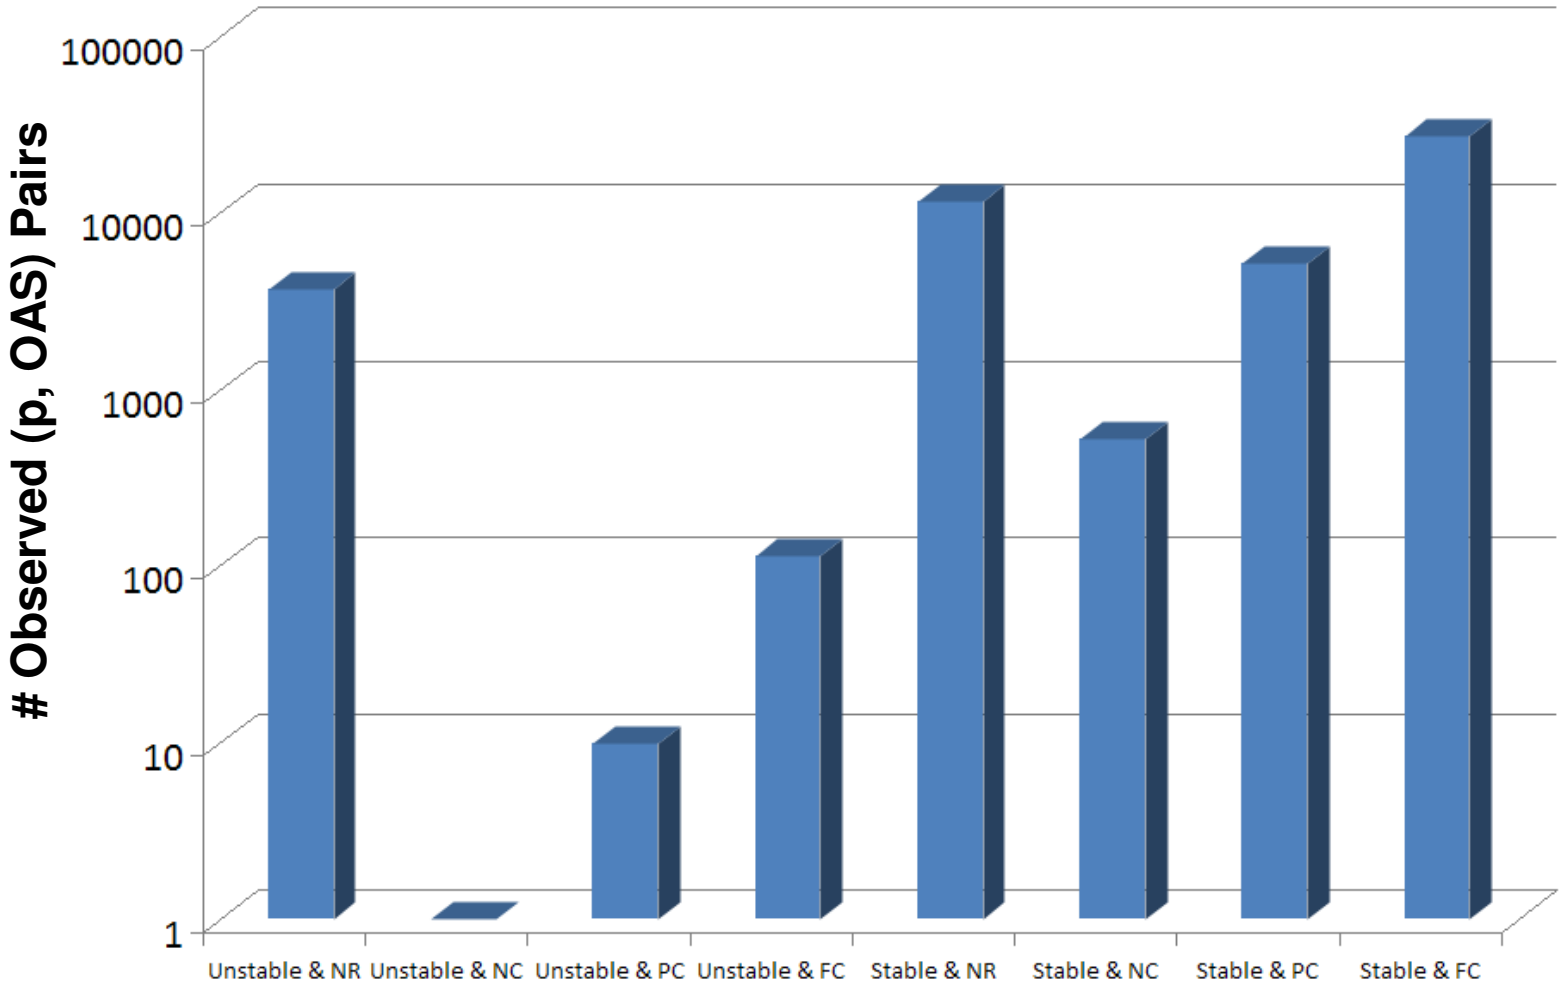

p = prefix; OAS = Origin AS; FC = Fully Consistent; PC = Partially Consistent; NC = Not Consistent; NR = Not Registered

Classification of Observed (p, OAS) Pairs According to Stability & Consistency Checks

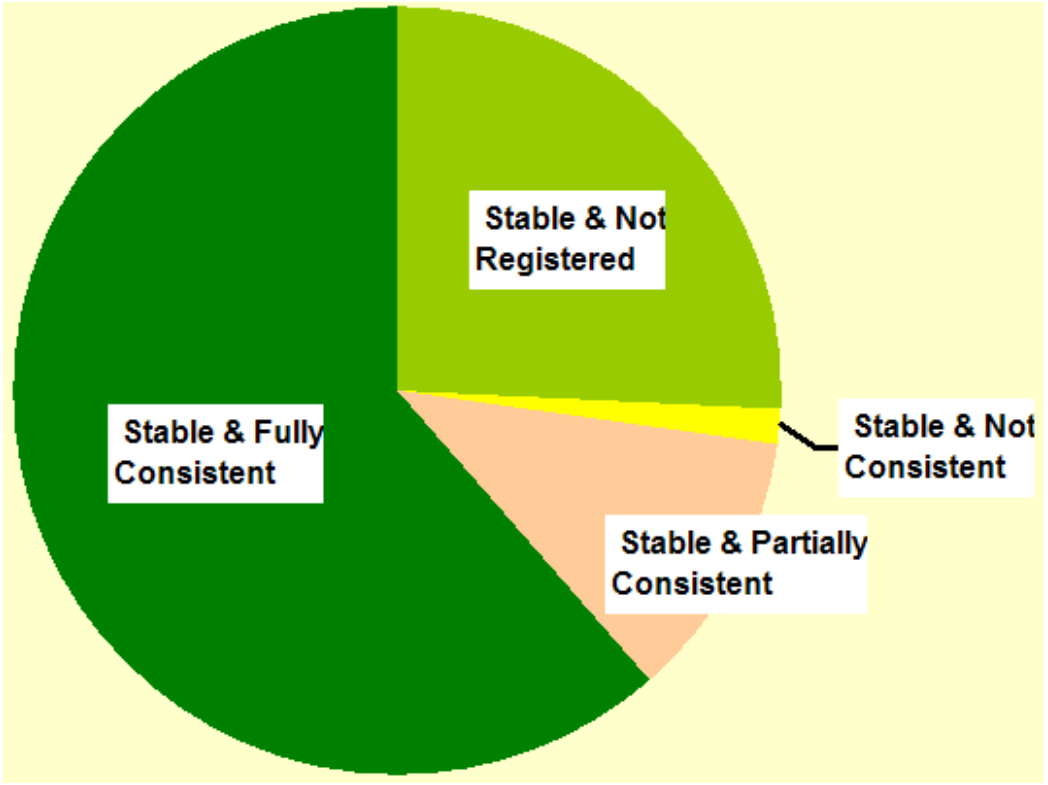

92% of (p, OAS) pairs are Stable

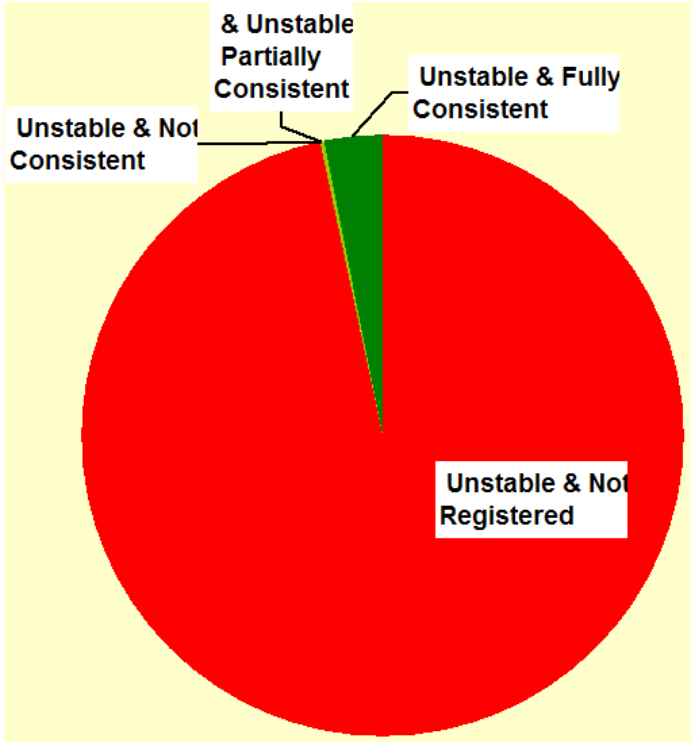

8 % of (p, OAS) pairs are Unstable

Heatmap Depicting Origin Validation for Prefixes in RIPE Region

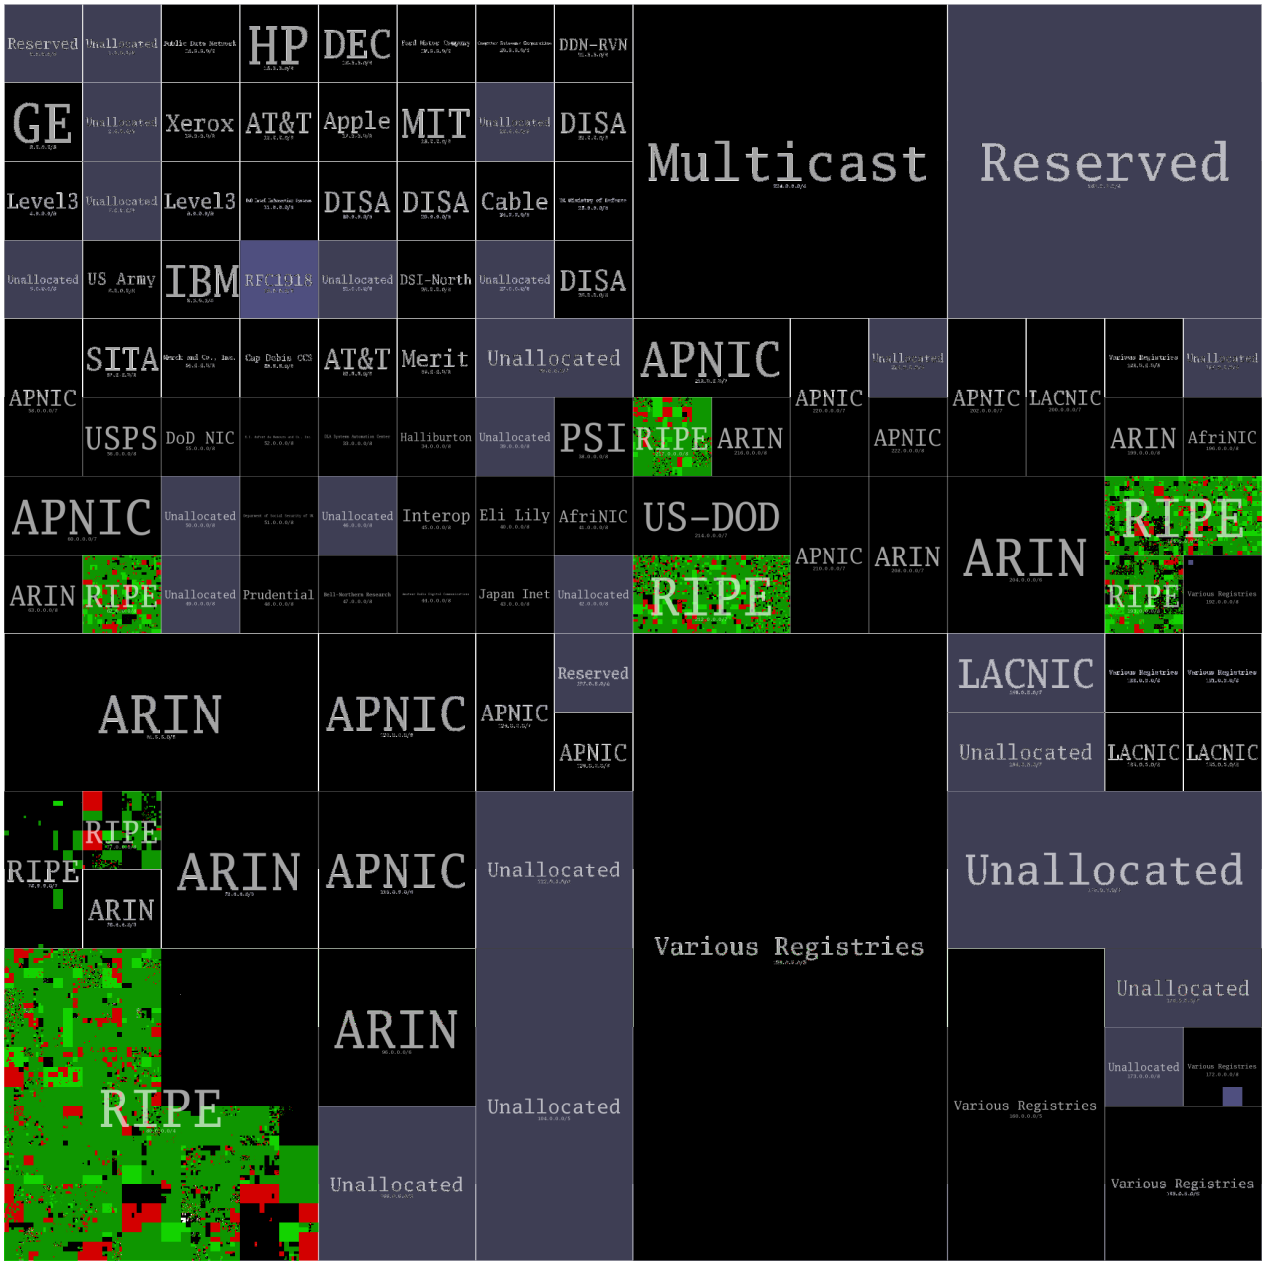

- **Registry-based Algorithm**
- **Trace-data (p, OAS) pairs are filtered based on RIPE NCC addresses**
- **So non-RIPE blocks are not scored**

Green: Good / FC

Light Green: Good / PC

Red: Suspicious

Black: Not found in trace data

Reference: http://maps.measurement-factory.com/software/ipv4-heatmap.1.html

Heatmap Depicting Origin Validation for Prefixes in RIPE Region

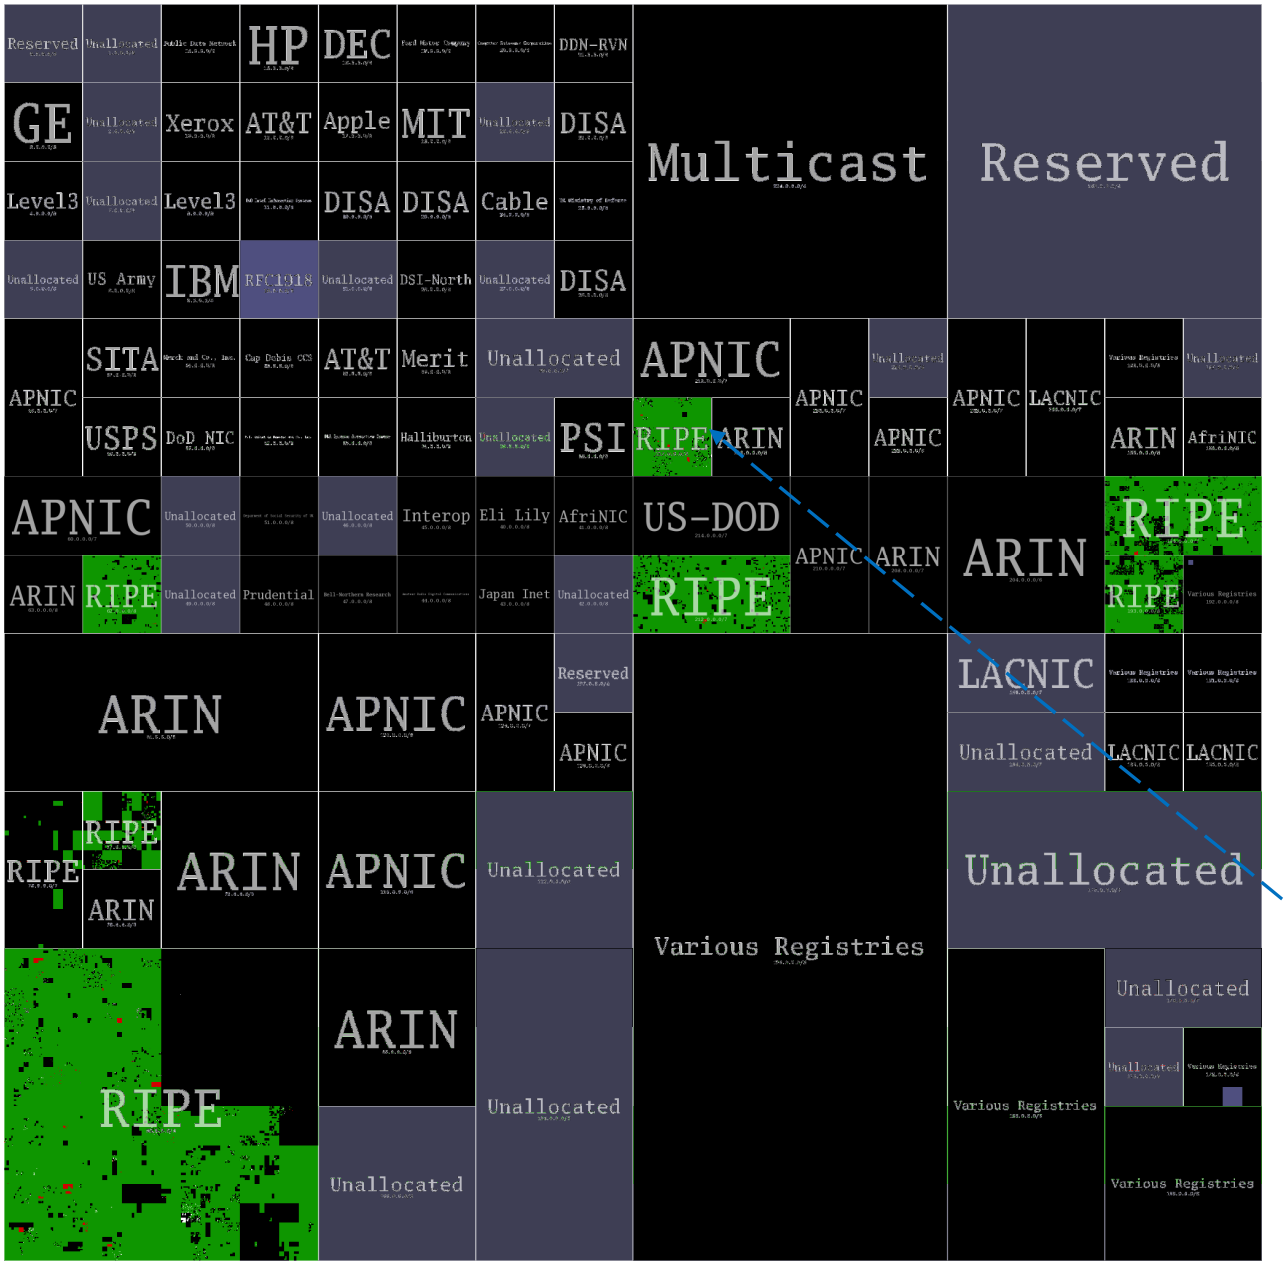

- **Trace-data based Algorithm**

Green: Good / FC

Light Green: Good / PC

Red: Suspicious

Black: Not found in trace data

We will zoom into this 217.0.0.0/8

Checking Origin AS : Comparison of Algorithms

- **Zooming into a RIPE 217.0.0.0/8 address block**

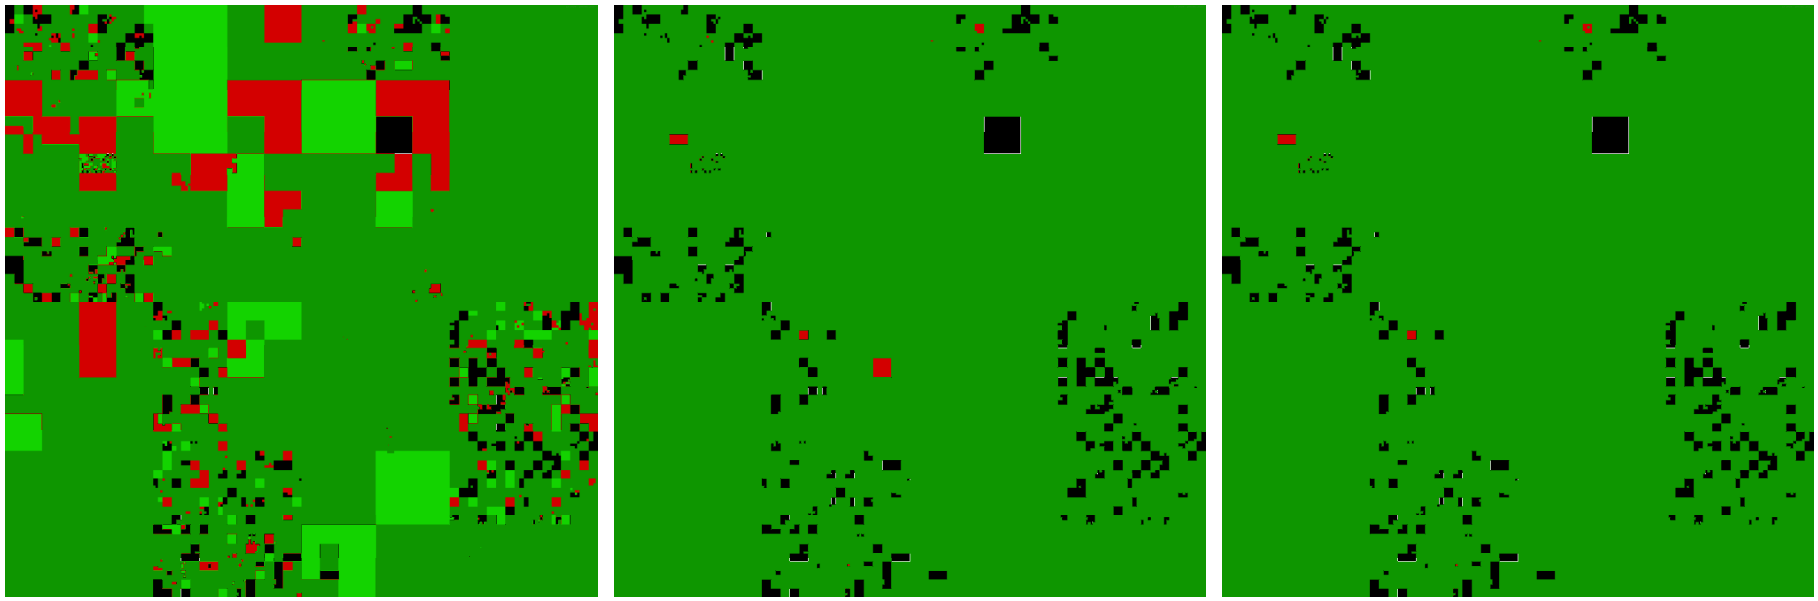

Registry-based Algorithm | **Trace-data based Algorithm** | **Enhanced Hybrid Algorithm**

Improvement in Anomaly Detection Algorithm
(Decreasing rate of false positives)

Green: Good / FC

Light Green: Good / PC

Red: Suspicious

Black: Not found in trace data

Checking Origin AS : Comparison of Algorithms

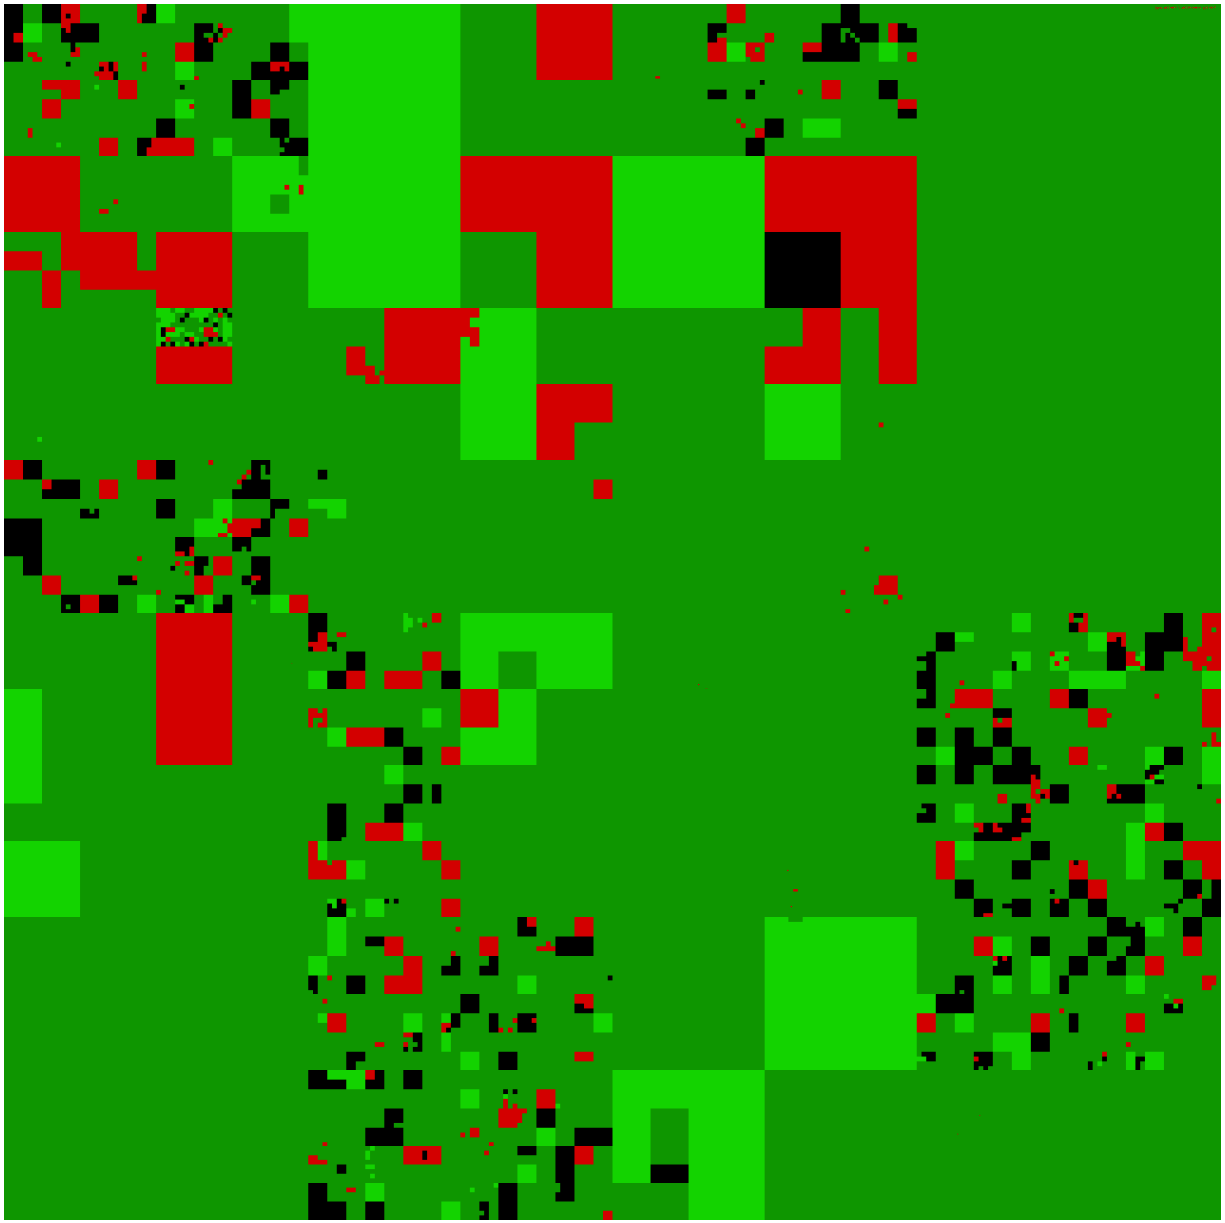

Registry-based Algorithm

Zooming into RIPE 217.0.0.0/8

Green: Good / FC

Light Green: Good / PC

Red: Suspicious

Black: Not found in trace data

Checking Origin AS : Comparison of Algorithms

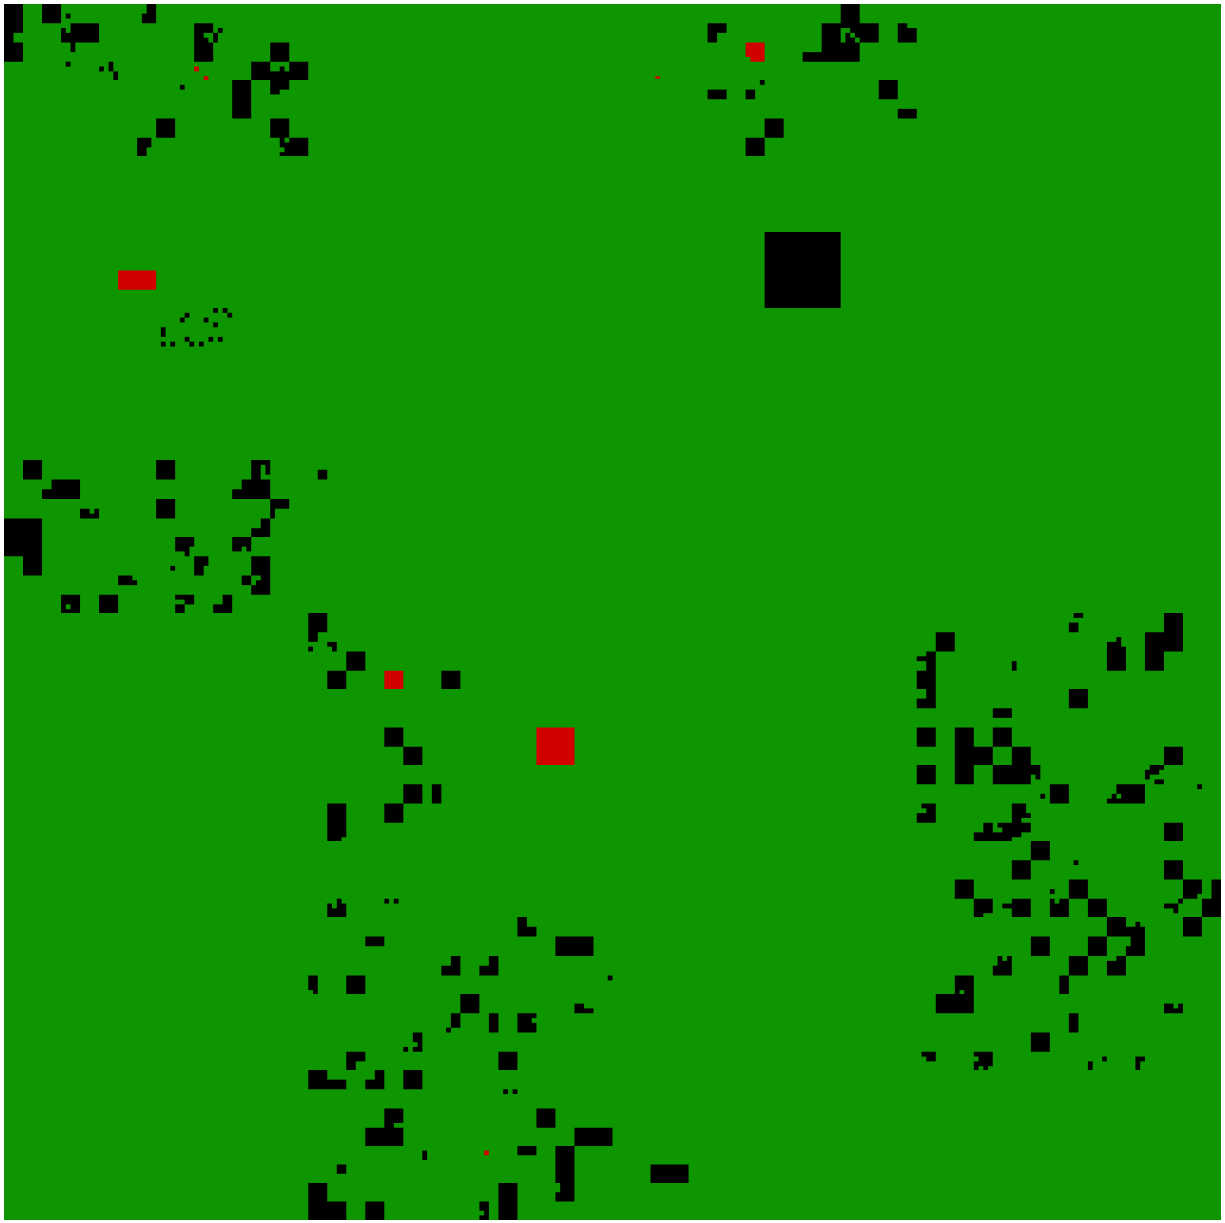

Trace-data based Algorithm

Zooming into RIPE 217.0.0.0/8

Green: Good / FC

Light Green: Good / PC

Red: Suspicious

Black: Not found in trace data

Checking Origin AS : Comparison of Algorithms

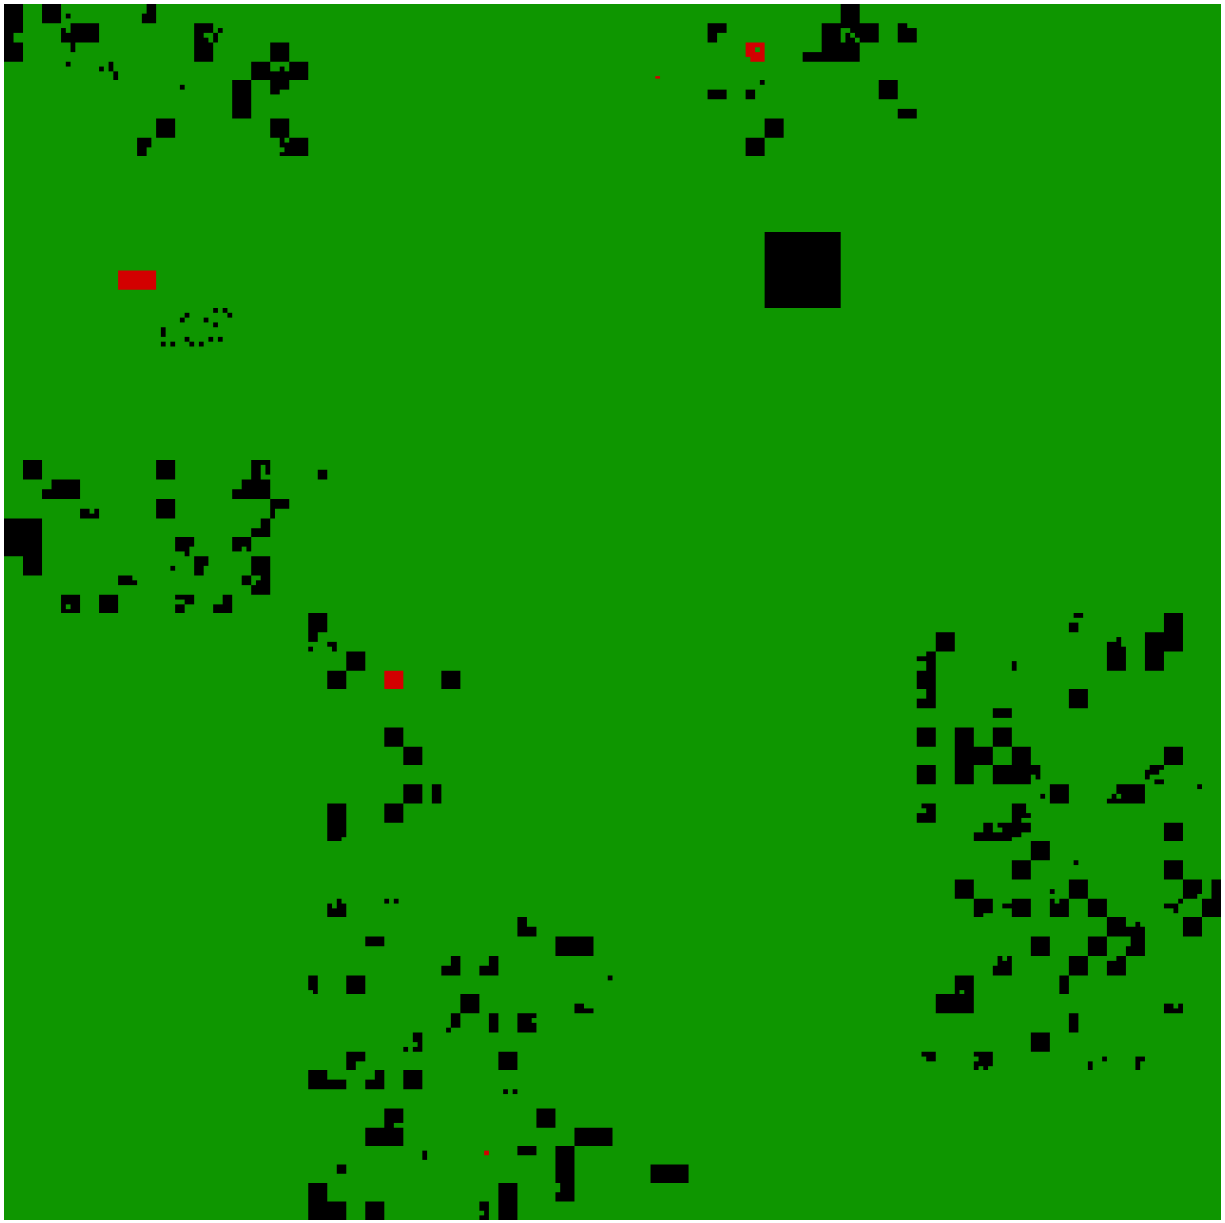

Enhanced Hybrid Algorithm

Zooming into RIPE 217.0.0.0/8

Green: Good / FC

Light Green: Good / PC

Red: Suspicious

Black: Not found in trace data

Comparative Performance of Algorithms

(RIPE NCC Filtered)

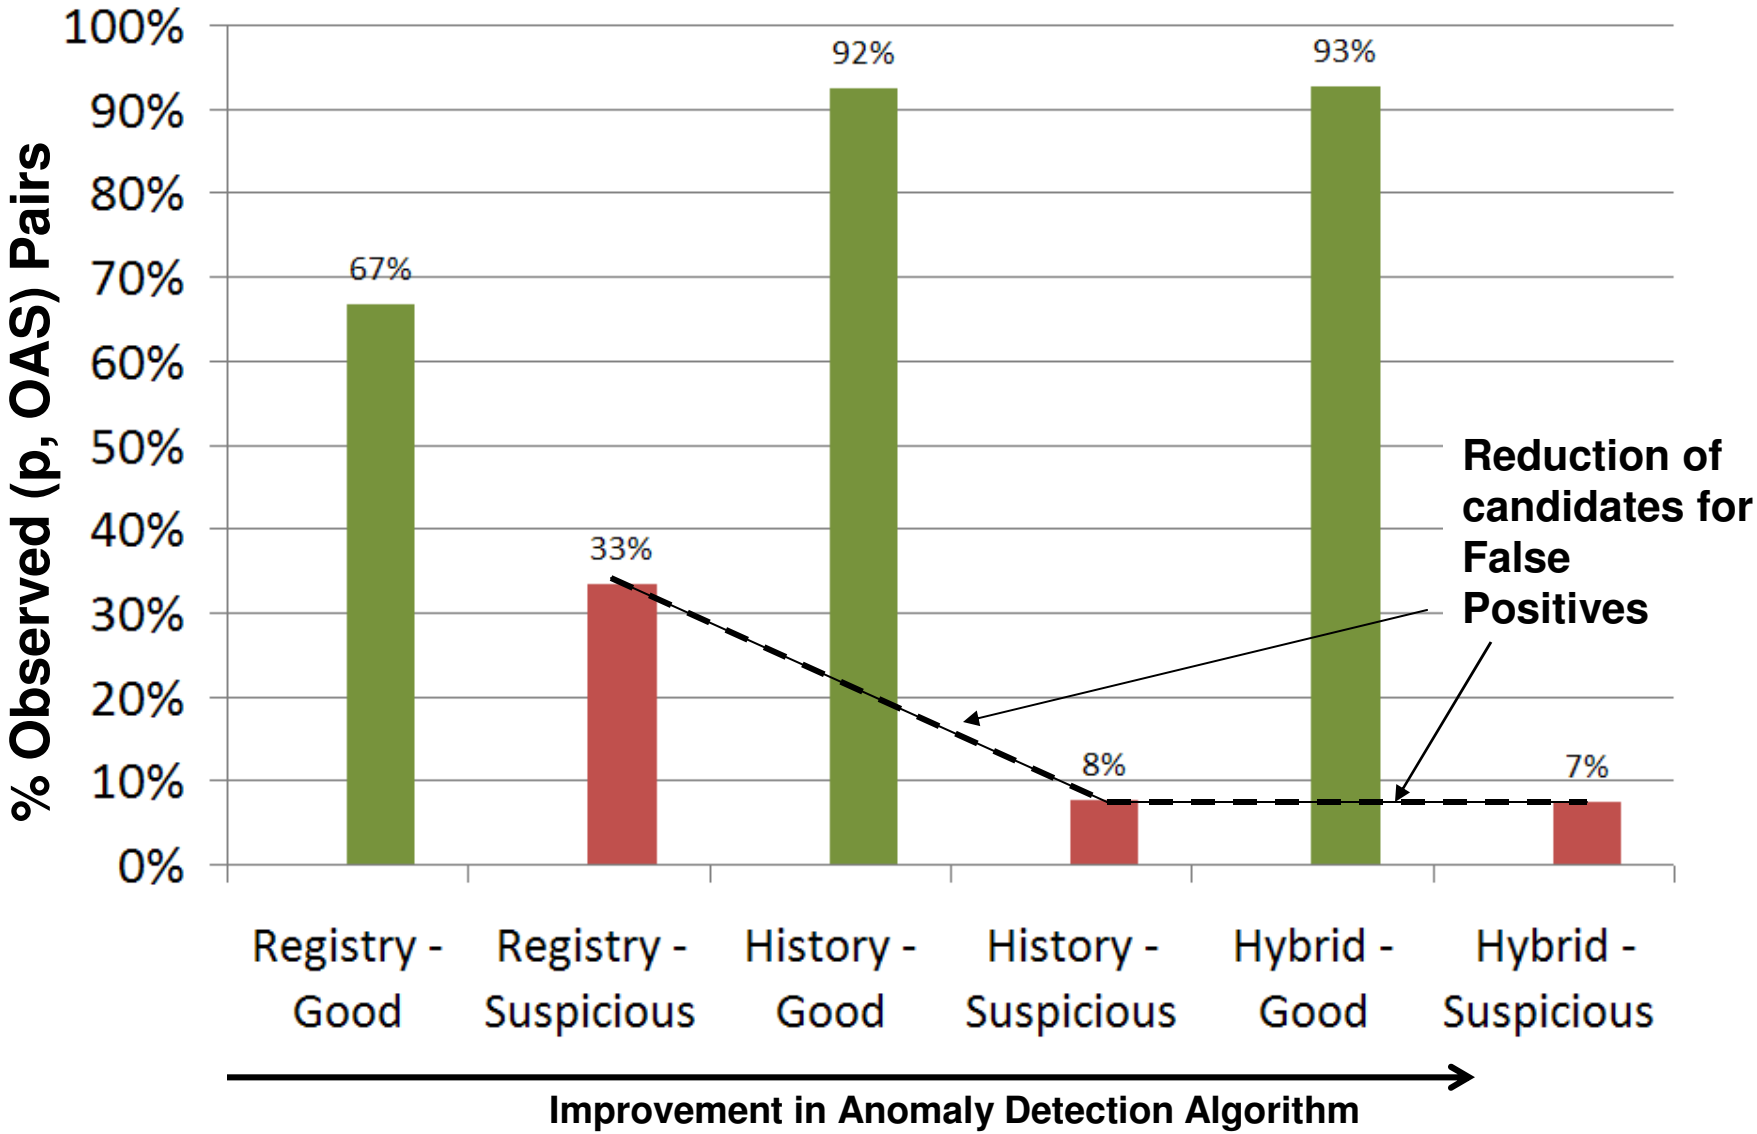

Prefixes with Multiple Origin ASes

# Origin ASes	# Prefixes
1	48972
2	503
3	9

For prefixes with two Origin ASes:

OAS1	OAS2	# Prefixes
FC + Stable	FC/PC + Unstable	4
PC + Stable	FC/PC + Unstable	1
NC + Stable	FC/PC + Unstable	0
NR + Stable	FC/PC + Unstable	2

- In some cases of prefixes with multiple Origin ASes, the primary path is stable (with or without consistency in the registry), while the secondary (failover) path is transient (unstable) but consistent in the registry

Conclusions and Planned Future Work

- Presented an overview and comparisons of BGP robustness and anomaly detection algorithms
- Several ***caveats*** apply in the reported results (To Do list)
  - Consideration of registry data from all regions
  - Reconciling related registry objects in different regional registries
  - Consideration of multiple trace-data collectors
- Work in progress – many more details being worked
- Further testing for robustness of the algorithms will be performed with extensive real and synthetic trace data.
- This will lead to numerical results for benchmarking the algorithms
- Help industry understand implications of proposals emerging from various ongoing R&D projects

Thank you!

Questions?

Updated slides and our other BGP Security papers/presentations:
http://www.antd.nist.gov/~ksriram/

Backup Slides

Details of Algorithm to Establish Stability of (p, Origin AS) in the Trace Data

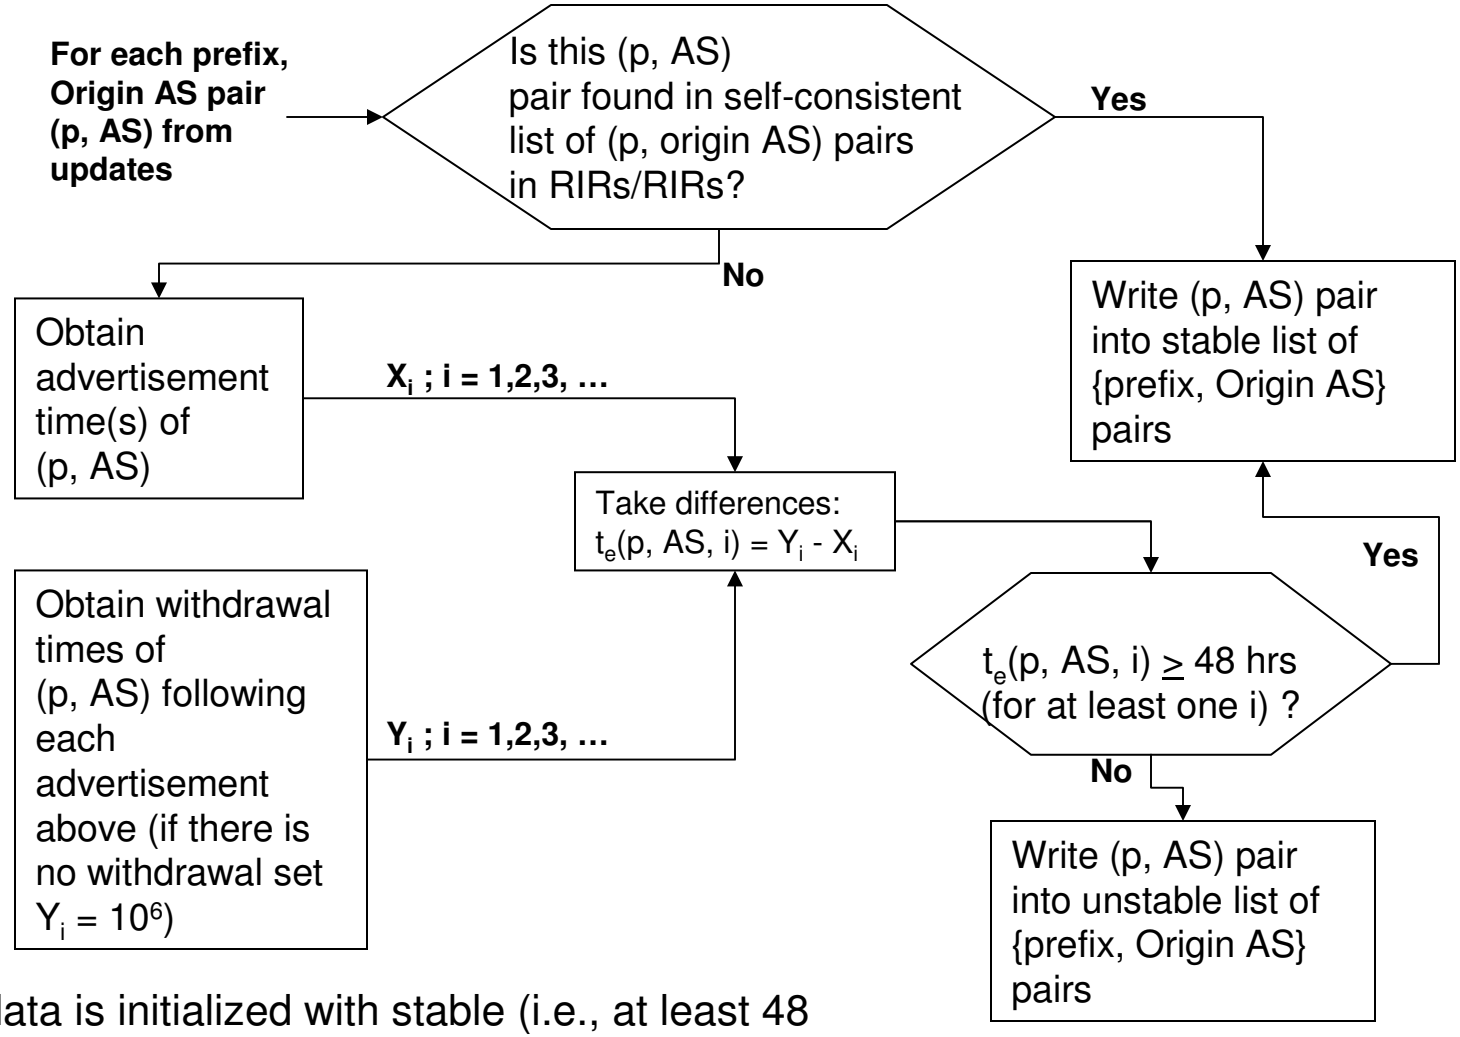

- Update data is initialized with stable (i.e., at least 48 hours) RIB entries

Trace-Data Based Algorithm: Differences Relative to PGBGP

- PGBGP considers a moving 10-day window of trace data
- We keep in our stable list any (p, OAS) pair that remained in the RIB for 48 hours or more at least once in our observation period
- The idea is that backup protection paths may be infrequently used
  - An AS may have served as the origin AS a few months ago during failover and is used again now
  - It is better to make that part of "stable" history if the (p, OAS) pair earlier remained in RIB for 48 hours or more
- We also augment the above with consideration of registry consistency checks in our enhanced hybrid algorithm

YouTube Hijack: Background Information

Prefix normally advertised by YouTube: 208.65.152.0/22 via AS 36561

Related (overlapping) prefixes seen historically and stayed stable for 48-hour or more:

Prefix	Origin AS	AS name	Time
208.65.152.0/22	AS 36561	YOUTUBE: YouTube, Inc.	02-20-08 15:43:50 (RIPE RIS) 02-20-08 15:37:46 (rrc02)

YouTube Hijack: Sequence of Events

Prefix normally advertised by YouTube: 208.65.152.0/22 via AS 36561

Date: 2/20/08 15:43:50	Normal announcement of 208.65.152.0/22 by AS 36561
15:37:46	rrc02: Prefix: 208.65.152.0/22, Origin: 36561, AS path: 14361 36561
Date: 2/24/08	
18:47:45	first evidence of hijacked route propagating in Asia, AS path 3491 17557 (208.65.153.0/24)
18:37:46	rrc02: Prefix: 208.65.153.0/24, Origin: 17557, AS path: 2497 3491 17557
18:49:00	most of the DFZ now carrying the bad route (and 93 ASNs)
18:49:30	all providers who will carry the hijacked route have it (total 97 ASNs)
20:07:25	YouTube, AS 36561 advertises the /24 that has been hijacked
20:07:25	rrc02: Prefix: 208.65.153.0/24, Origin: 36561, AS path:19089 3549 36561
20:08:30	a total of 40 some-odd providers have stopped using the hijacked route

Notes: rrc02 update data (yellow rows) is from TERRAIN database

Event timeline (white rows) obtained from Martin A. Brown's blog at Renesys:
http://www.renesys.com/blog/2008/02/pakistan_hijacks_youtube_1.shtml

Continued on next page …

YouTube Hijack: Sequence of Events (Contd.)

Prefix normally advertised by YouTube: 208.65.152.0/22 via AS 36561

Date: 2/24/08	
20:18:43	and now, two more specific /25 routes are first seen from 36561
20:18:43 20:18:43	rrc02: Prefix: 208.65.153.0/25, Origin: 36561, AS path:19089 3549 36561 rrc02: Prefix: 208.65.153.128/25, Origin: 36561, AS path: 19089 3549 36561
20:19:37	25 more providers prefer the /25 routes from 36561
20:50:59	evidence of attempted prepending, AS path was 3491 17557 17557
20:59:39	hijacked prefix is withdrawn by 3491, who disconnect 17557

Notes: rrc02 update data (yellow rows) is from TERRAIN database

Event timeline (white rows) obtained from Martin A. Brown's blog at Renesys:
http://www.renesys.com/blog/2008/02/pakistan_hijacks_youtube_1.shtml

How Effective in a YouTube Like Incident:

Detecting and Alerting the Attack by Pakistan Telecom

		Results of checks included in each approach			
		Registry-based approach	Trace-data based approach "PGBGP"	Simple Hybrid	Enhanced Hybrid
Q1.	Is prefix registered (same or less specific)?	**No**		**No**	**No**
Q2.	Is there a route registered (with same or less specific prefix and origin AS)?	**No**		**No**	**No**
Q3.	Is announced (p, origin AS) fully consistent with corresponding registry objects in RIR/IRR?	**No**		**No**	**No**
Q4.	Is announced (p, origin AS) partially consistent with corresponding registry objects RIR/IRR?	**No**		**No**	**No**
Q5.	Was (p, origin AS) seen in RIB in the last *h* (= 10) days? (Also, if it was suspicious, did it remain in RIB beyond the suspicious period of *s* (= 24) hours?)		**No**	**No**	
Q6.	Would a less specific prefix with the same origin AS pass the test in Q5?		**No**	**No**	
Q7.	Was prefix previously announced by the same origin AS and remained stably (48 hrs or more) in the RIB over the observation period (*d* months)?				**No**
Q8.	Would a less specific prefix with the same origin AS pass the test in Q7?				**No**
Q9.	Is the peering rank of the origin AS high or medium?				**No**

How Effective in a YouTube Like Incident:

Detecting and Allowing Recovery Using Sub-prefixes by YouTube

		Results of checks included in each approach			
	Checks/Questions	Registry-based approach	Trace-data based approach "PGBGP"	Simple Hybrid	Enhanced Hybrid
Q1.	Is prefix registered (same or less specific)?	**Yes**		**Yes**	**Yes**
Q2.	Is there a route registered (with same or less specific prefix and origin AS)?	**Yes**		**Yes**	**Yes**
Q3.	Is announced (p, origin AS) fully consistent with corresponding registry objects in RIR/IRR?	**Yes**		**Yes**	**Yes**
Q4.	Is announced (p, origin AS) partially consistent with corresponding registry objects RIR/IRR?	**Yes**		**Yes**	**Yes**
Q5.	Was (p, origin AS) seen in RIB in the last *h* (= 10) days? (Also, if it was suspicious, did it remain in RIB beyond the suspicious period of *s* (= 24) hours?)		**No**	**No**	
Q6.	Would a less specific prefix with the same origin AS pass the test in Q5?		**Yes**	**Yes**	
Q7.	Was prefix previously announced by the same origin AS and remained stably (48 hrs or more) in the RIB over the observation period (*d* months)?				**No**
Q8.	Would a less specific prefix with the same origin AS pass the test in Q7?				**Yes**
Q9.	Is the peering rank of the origin AS high or medium?				**Yes**

YouTube Hijack: Actions by Different Algorithms

Time	Event	Registry -based	PHAS	PGBGP	Enhanced Hybrid Algorithm
Date: 2/20/08 15:43:50Z	Normal /22 Re-Advertisement	No alert	No alert	Propagate update	Propagate update
Date: 2/24/08 15:37:46	Hijack attempt with /24 subprefix	Alert	Alert: new origin	Quarantine update	Quarantine update
18:37:46	Recovery attempt with /24 subprefix	No alert	Alert: Notify subprefix	Propagate update	Propagate update
20:07:25	Recovery attempt with /25 subprefix	No alert	Alert: Notify subprefix	Propagate update	Propagate update

- The proposed enhanced hybrid algorithms would effectively deal with certain special situations that did not manifest in this set of events.

Evaluation of BGP Anomaly Detection and Robustness Algorithms

Abstract:
We present an evaluation methodology for comparison of existing and proposed new algorithms for BGP anomaly detection and robustness. A variety of algorithms and alert tools have been proposed and/or prototyped recently. They differ in the anomaly situations which they attempt to alert or mitigate, and also in the type(s) of data they use. Some are based on registry data from RIRs/IRRS (e.g. Nemecis) and others (PHAS, PGBGP) are driven by BGP trace data. The trace data is obtained from RIPE-RIS, Route-views, or a BGP speaker where the algorithm operates. We propose a new algorithm that combines the use of both registry and trace data, and also makes some key improvements over existing algorithms. We have built an evaluation platform called TERRAIN (Testing and Evaluation of Routing Robustness in Assurable Inter-domain Networking) on which these algorithms can be tested and empirically compared based on real and/or synthetically incorporated anomalies in BGP updates. We will present a variety of preliminary results providing interesting insights into the comparative utility and performance of the various BGP robustness algorithms. Our objective is to share these early insights and invite feedback from the community to refine the TERRAIN evaluation tool to generate further useful results in the future.